

Introduction to Database Systems

Relational Algebra

Werner Nutt

1

Motivation

- We know how to store data ...
- How can we retrieve (interesting) data?
- We need a query language
 - *declarative* (to allow for abstraction)
 - *optimisable* (→ less expressive than a programming language, not Turing-complete)
 - relations as input and output

E.F. Codd (1970): Relational Algebra

2

Characteristics of an Algebra

- Expressions
 - Are constructed with operators from atomic operands (constants, variables,)
 - can be evaluated
- Expressions can be equivalent
 - ...if they return the same result for all values of the variables
 - Equivalence gives rise to identities between (schemas of) expressions
- The value of an expression is independent of its context
 - e.g., $5 + 3$ has the same value, no matter whether it occurs as
$$10 - (5 + 3) \quad \text{or} \quad 4 \cdot (5 + 3)$$
 - Consequence: subexpressions can be replaced by equivalent expressions without changing the meaning of the entire expression

3

Example: Algebra of Arithmetic

- Atomic expressions:
numbers and variables
- Operators: $+, -, \cdot, :$
- Identities:
$$x + y = y + x$$
$$x \cdot (y + z) = x \cdot y + x \cdot z$$
$$\dots$$

4

Relational Algebra: Principles

Atoms are relations

Operators are defined for arbitrary instances of a relation

Two results have to be defined for each operator

1. result schema
(depending on the schemas of the argument relations)
2. result instance
(depending on the instances of the arguments)

- *“Equivalent” to SQL query language*
- *Relational Algebra concepts reappear in SQL*
- *Used inside a DBMS, to express query plans*

5

Classification of Relational Algebra Operators

- Set theoretic operators
union “ \cup ”, intersection “ \cap ”, difference “ \setminus ”
- Renaming operator ρ
- Removal operators
projection π , selection σ
- Combination operators
Cartesian product “ \times ”, joins “ \bowtie ”
- Extended operators
duplicate elimination, grouping, aggregation, sorting,
outer joins, etc.

6

Set Theoretic Operators

Observations:

- Instances of relations are sets
→ we can form unions, intersections, and differences
- Results of algebra operations must be relations, i.e., results must have a schema

Would it make sense to apply these operators to bags (= multisets)?

Hence:

- Set theoretic algebra operators can only be applied to relations with identical attributes, i.e.,
 - same *number* of attributes
 - same *names*
 - same *domains*

7

Union

CS-Student

Studno	Name	Year
s1	Egger	5
s3	Rossi	4
s4	Maurer	2

Master-Student

Studno	Name	Year
s1	Egger	5
s2	Neri	5
s3	Rossi	4

CS-Student \cup Master-Student

Studno	Name	Year
s1	Egger	5
s2	Neri	5
s3	Rossi	4
s4	Maurer	2

*Note: relations are sets
→ no duplicates
are eliminated*

8

Intersection

CS-Student

Studno	Name	Year
s1	Egger	5
s3	Rossi	4
s4	Maurer	2

Master-Student

Studno	Name	Year
s1	Egger	5
s2	Neri	5
s3	Rossi	4

CS-Student \cap Master-Student

Studno	Name	Year
s1	Egger	5
s3	Rossi	4

9

Difference

CS-Student

Studno	Name	Year
s1	Egger	5
s3	Rossi	4
s4	Maurer	2

Master-Student

Studno	Name	Year
s1	Egger	5
s2	Neri	5
s3	Rossi	4

CS-Student \setminus Master-Student

Studno	Name	Year
s4	Maurer	2

10

Not Every Union That Makes Sense is Possible

Father-Child

Father	Child
Adam	Abel
Adam	Cain
Abraham	Isaac

Mother-Child

Mother	Child
Eve	Abel
Eve	Seth
Sara	Isaac

Father-Child \cup Mother-Child

??

11

Renaming

- The renaming operator ρ changes the name of one or more attributes
- It changes the schema, but not the instance of a relation

Father-Child

Father	Child
Adam	Abel
Adam	Cain
Abraham	Isaac

$\rho_{\text{Parent} \leftarrow \text{Father}}(\text{Father-Child})$

Parent	Child
Adam	Abel
Adam	Cain
Abraham	Isaac

12

Father-Child

Father	Child
Adam	Abel
Adam	Cain
Abraham	Isaac

$\rho_{\text{Parent} \leftarrow \text{Father}}$ (Father-Child)

Parent	Child
Adam	Abel
Adam	Cain
Abraham	Isaac

Mother-Child

Mother	Child
Eve	Abel
Eve	Seth
Sara	Isaac

$\rho_{\text{Parent} \leftarrow \text{Mother}}$ (Mother-Child)

Parent	Child
Eve	Abel
Eve	Seth
Sara	Isaac

$\rho_{\text{Parent} \leftarrow \text{Father}}$ (Father-Child)

Parent	Child
Adam	Abel
Adam	Cain
Abraham	Isaac

$\rho_{\text{Parent} \leftarrow \text{Father}}$ (Father-Child)

∪

$\rho_{\text{Parent} \leftarrow \text{Mother}}$ (Mother-Child)

$\rho_{\text{Parent} \leftarrow \text{Mother}}$ (Mother-Child)

Parent	Child
Eve	Abel
Eve	Seth
Sara	Isaac

Parent	Child
Adam	Abel
Adam	Cain
Abraham	Isaac
Eve	Abel
Eve	Seth
Sara	Isaac

Projection and Selection

Two “orthogonal” operators

- Selection:
 - horizontal decomposition
- Projection:
 - vertical decomposition



15

Projection

General form:

$$\pi_{A_1, \dots, A_k}(R)$$

where R is a relation and A_1, \dots, A_k are attributes of R .

Result:

- Schema: (A_1, \dots, A_k)
- Instance: the set of all subtuples $t[A_1, \dots, A_k]$ where $t \in R$

Intuition: Deletes all attributes that are not in projection list

In general, needs to *eliminate duplicates* *projection without this!*
... but not if A_1, \dots, A_k comprises a key (why?)

16

Projection: Example

STUDENT

studno	name	hons	tutor	year
s1	jones	ca	bush	2
s2	brown	cis	kahn	2
s3	smith	cs	goble	2
s4	bloggs	ca	goble	1
s5	jones	cs	zobel	1
s6	peters	ca	kahn	3

$\pi_{\text{tutor}}(\text{STUDENT}) =$

tutor
bush
kahn
goble
zobel

Note: result relations don't have a name

17

Selection

General form:

$$\sigma_C(R)$$

with a relation R and a condition C on the attributes of R.

Result:

- Schema: the schema of R
- Instance: the set of all $t \in R$ that satisfy C

Intuition: Filters out all tuples that do not satisfy C

No need to eliminate duplicates (Why?)

18

Selection: Example

STUDENT

studno	name	hons	tutor	year
s1	jones	ca	bush	2
s2	brown	cis	kahn	2
s3	smith	cs	goble	2
s4	bloggs	ca	goble	1
s5	jones	cs	zobel	1
s6	peters	ca	kahn	3

$\sigma_{\text{name='bloggs'}}(\text{STUDENT}) =$

studno	name	hons	tutor	year
s4	bloggs	ca	goble	1

Selection Conditions

Elementary conditions:

$\langle \text{attr} \rangle \text{ op } \langle \text{val} \rangle$ or $\langle \text{attr} \rangle \text{ op } \langle \text{attr} \rangle$ or $\langle \text{expr} \rangle \text{ op } \langle \text{expr} \rangle$

where op is "=", "<", "≤", (on numbers and strings)
"LIKE" (for string comparisons),...

No specific set of elementary conditions is built into rel alg ...

Example: age ≤ 24, phone LIKE '0039%',
salary + commission ≥ 24 000

Combined conditions (using Boolean connectives):

C1 and C2 or C1 or C2 or not C

Operators Can Be Nested

Who is the tutor of the student named “Bloggs”?

STUDENT				
studno	name	hons	tutor	year
s1	jones	ca	bush	2
s2	brown	cis	kahn	2
s3	smith	cs	goble	2
s4	bloggs	ca	goble	1
s5	jones	cs	zobel	1
s6	peters	ca	kahn	3

$$\pi_{\text{tutor}} (\sigma_{\text{name}='bloggs'} (\text{STUDENT}))$$
$$= \begin{array}{|c|} \hline \text{tutor} \\ \hline \text{goble} \\ \hline \end{array}$$

21

Operator Applicability

Not every operator can be applied to every relation:

- Projection: π_{A_1, \dots, A_k} is applicable to R if
R has attributes with the names A_1, \dots, A_k
- Selection: σ_C is applicable to R if
all attributes mentioned in C
appear as attributes of R
and have the right types

22

Identities for Selection and Projection

For all conditions C1, C2 and relations R we have:

- $\sigma_{C1}(\sigma_{C2}(R)) = \sigma_{C2}(\sigma_{C1}(R)) = \sigma_{C1 \text{ and } C2}(R)$

What about

- $\pi_{A1, \dots, Am}(\pi_{B1, \dots, Bn}(R)) = \pi_{B1, \dots, Bn}(\pi_{A1, \dots, Am}(R))$?

And what about

- $\pi_{A1, \dots, Am}(\sigma_C(R)) = \sigma_C(\pi_{A1, \dots, Am}(R))$?

Any ideas for more identities?

23

Selection Conditions and “null”

Does the following identity hold:

$$\text{Student} = \sigma_{\text{year} \leq 3}(\text{Student}) \cup \sigma_{\text{year} > 3}(\text{Student}) ?$$

What if Student contains a tuple t with t[year] = null ?

Convention: Only comparisons with non-null values are TRUE or FALSE.
Conceptually, comparisons involving null yield a value UNKNOWN.

To test, whether a value is null or not null, there are two conditions:

<attr> IS NULL
<attr> IS NOT NULL

24

Identities with “null”

Thus, the following identities hold:

$$\begin{aligned}\text{Student} &= \sigma_{\text{year} \leq 3}(\text{Student}) \cup \sigma_{\text{year} > 3}(\text{Student}) \cup \sigma_{\text{year IS NULL}}(\text{Student}) \\ &= \sigma_{\text{year} \leq 3 \text{ OR } \text{year} > 3 \text{ OR } \text{year IS NULL}}(\text{Student})\end{aligned}$$

25

Exercises

Write relational algebra queries that retrieve:

1. All staff members that lecture or tutor
2. All staff members that lecture and tutor
3. All staff members that lecture, but don't tutor

STUDENT

studno	name	hons	tutor	year
s1	jones	ca	bush	2
s2	brown	cis	kahn	2
s3	smith	cs	goble	2
s4	bloggs	ca	goble	1
s5	jones	cs	zobel	1
s6	peters	ca	kahn	3

TEACH

courseno	lecturer
cs250	lindsey
cs250	capon
cs260	kahn
cs260	bush
cs270	zobel
cs270	woods
cs280	capon

26

Cartesian Product

General form: $R \times S$

where R and S are arbitrary relations

Result:

- Schema: $(A_1, \dots, A_m, B_1, \dots, B_n)$, if (A_1, \dots, A_m) is the schema of R and (B_1, \dots, B_n) is the schema of S.

(If A is an attribute of both, R and S, then $R \times S$ contains the *disambiguated attributes* R.A and S.A.)

- Instance: the set of all *concatenated tuples*

(t, s)

where $t \in R$ and $s \in S$

27

Cartesian Product: Student \times Course

STUDENT

studno	name
s1	jones
s2	brown
s6	peters

COURSE

courseno	subject	equip
cs250	prog	sun
cs150	prog	sun
cs390	specs	sun

STUDENT \times COURSE

studno	name	courseno	subject	equip
s1	jones	cs250	prog	sun
s1	jones	cs150	prog	sun
s1	jones	cs390	specs	sun
s2	brown	cs250	prog	sun
s2	brown	cs150	prog	sun
s2	brown	cs390	specs	sun
s6	peters	cs250	prog	sun
s6	peters	cs150	prog	sun
s6	peters	cs390	specs	sun

28

Cartesian Product: Student \times Staff

STUDENT

studno	name	hons	tutor	year
s1	jones	ca	bush	2
s2	brown	cis	kahn	2
s3	smith	cs	goble	2
s4	bloggs	ca	goble	1
s5	jones	cs	zobel	1
s6	peters	ca	kahn	3

STAFF

lecturer	roomno
kahn	IT206
bush	2.26
goble	2.82
zobel	2.34
watson	IT212
woods	IT204
capon	A14
lindsey	2.10
barringer	2.125

*What's
the point
of this?*

studno	name	hons	tutor	year	lecturer	roomno
s1	jones	ca	bush	2	kahn	IT206
s1	jones	ca	bush	2	bush	2.26
s1	jones	ca	bush	2	goble	2.82
s1	jones	ca	bush	2	zobel	2.34
s1	jones	ca	bush	2	watson	IT212
s1	jones	ca	bush	2	woods	IT204
s1	jones	ca	bush	2	capon	A14
s1	jones	ca	bush	2	lindsey	2.10
s1	jones	ca	bush	2	barringer	2.125
s2	brown	cis	kahn	2	kahn	IT206
s2	brown	cis	kahn	2	bush	2.26
s2	brown	cis	kahn	2	goble	2.82
s2	brown	cis	kahn	2	zobel	2.34
s2	brown	cis	kahn	2	watson	IT212
s2	brown	cis	kahn	2	woods	IT204
s2	brown	cis	kahn	2	capon	A14
s2	brown	cis	kahn	2	lindsey	2.10
s2	brown	cis	kahn	2	barringer	2.125
s3	smith	cs	goble	2	kahn	IT206
s3	smith	cs	goble	2	bush	2.26
s3	smith	cs	goble	2	goble	2.82
s3	smith	cs	goble	2	zobel	2.34
s3	smith	cs	goble	2	watson	IT212
s3	smith	cs	goble	2	woods	IT204
s3	smith	cs	goble	2	capon	A14
s3	smith	cs	goble	2	lindsey	2.10
s3	smith	cs	goble	2	barringer	2.125
s4	bloggs	ca	goble	1	kahn	IT206

...

29

“Where is the Tutor of Bloggs?”

To answer the query

“For each student, identified by name and student number, return the name of the tutor and their office number”

we have to

- combine tuples from Student and Staff
- that satisfy “Student.tutor=Staff.lecturer”
- and keep the attributes studno, name, tutor and lecturer.

In relational algebra:

$$\pi_{\text{studno,name,lecturer,roomno}}(\sigma_{\text{tutor=lecturer}}(\text{Student} \times \text{Staff}))$$

The part $\sigma_{\text{tutor=lecturer}}(\text{Student} \times \text{Staff})$ is a “join”.

30

Example: Student Marks in Courses

STUDENT

studno	name	hons	tutor	year
s1	jones	ca	bush	2
s2	brown	cis	kahn	2
s3	smith	cs	goble	2
s4	bloggs	ca	goble	1
s5	jones	cs	zobel	1
s6	peters	ca	kahn	3

*“For each student,
show the courses in which
they are enrolled
and their marks!”*

ENROL

stud no	course no	lab mark	exam mark
s1	cs250	65	52
s1	cs260	80	75
s1	cs270	47	34
s2	cs250	67	55
s2	cs270	65	71
s3	cs270	49	50
s4	cs280	50	51
s5	cs250	0	3
s6	cs250	2	7

First, do

$$R \leftarrow \sigma_{\text{Student.studno} = \text{Enrol.studno}} (\text{Student} \times \text{Enrol}),$$

then

$$\text{Result} \leftarrow \pi_{\text{studno, name, ... , exam_mark}}(R)$$

31

Natural Join

Suppose: R, S are two relations

with attributes A_1, \dots, A_m and B_1, \dots, B_n , resp.
and with common attributes D_1, \dots, D_k

The natural join of R and S is a relation that has as

Schema: all attributes occurring in R or S,
where common attributes occur only once,
i.e., the set of attributes $\mathbf{Attr} = \{A_1, \dots, A_m, B_1, \dots, B_n\}$

Instance: all tuples t over the attributes \mathbf{Attr} such that
 $t[A_1, \dots, A_m] \in R$ and $t[B_1, \dots, B_n] \in S$

Notation: $R \bowtie S$

32

Natural Join is a Derived Operation

The natural join of R and S can be written using

- Cartesian Product
- Selection
- Projection

$$R \bowtie S = \pi_{\text{Attr}}(\sigma_{R.D1=S.D1 \text{ and } \dots \text{ and } R.Dk=S.Dk} (R \times S))$$

33

θ -Joins (read, Theta-Joins), Equi-Joins

Most general form of join

- First, form Cartesian product of R and S
- Then, filter $R \times S$ with operators (abstractly written “ θ ”) relating attributes of R and S

Notation:

$$R \bowtie_C S = \sigma_C (R \times S)$$

Special case: If C is a conjunction of equalities, i.e.,

$$C = R.A1=S.B1 \text{ and } \dots \text{ and } R.Ai=S.Bi$$

then the θ -Join with condition C is called an equi-join.

Example:

$$\sigma_{\text{tutor=lecturer}}(\text{Student} \times \text{Staff}) = \text{Student} \bowtie_{\text{tutor=lecturer}} \text{Staff}$$

34

STUDENT

studno	name	hons	tutor	year
s1	jones	ca	bush	2
s2	brown	cis	kahn	2
s3	smith	cs	goble	2
s4	bloggs	ca	goble	1
s5	jones	cs	zobel	1
s6	peters	ca	kahn	3

STAFF

lecturer	roomno
kahn	IT206
bush	2.26
goble	2.82
zobel	2.34
watson	IT212
woods	IT204
capon	A14
lindsey	2.10
barringer	2.125

Student \bowtie $_{\text{tutor=lecturer}}$ Staff $(= \sigma_{\text{tutor=lecturer}}(\text{Student} \times \text{Staff}))$

stud no	name	hons	tutor	year	lecturer	roomno
s1	jones	ca	bush	2	bush	2.26
s2	brown	cis	kahn	2	kahn	IT206
s3	smith	cs	goble	2	goble	2.82
s4	bloggs	ca	goble	1	goble	2.82
s5	jones	cs	zobel	1	zobel	2.34
s6	peters	ca	kahn	3	kahn	IT206

35

Self Joins

For some queries we have to combine data coming from a single relation.

“Give all pairings of lecturers and appraisers, including their room numbers!”

We need two identical versions of the STAFF relation.

Question: How can we distinguish between the versions and their attributes?

Idea:

- Introduce temporary relations with new names
- Disambiguate attributes by prefixing them with the relation names

LEC \leftarrow STAFF,

APP \leftarrow STAFF

36

Self Joins (cntd.)

*“Give all pairings of lecturers and appraisers,
including their room numbers!”*

$R \leftarrow \pi_{\text{LEC.lecturer,LEC.roomno, (LEC} \bowtie \text{LEC.appraiser=APP.lecturer APP)}$
 $\text{LEC.appraiser,APP.roomno}$

Result $\leftarrow \rho_{\text{lecturer} \leftarrow \text{LEC.lecturer,roomno} \leftarrow \text{LEC.roomno, (R)}$
 $\text{appraiser} \leftarrow \text{LEC.appraiser,aproom} \leftarrow \text{APP.roomno}$

STAFF

lecturer	roomno	appraiser
kahn	IT206	watson
bush	2.26	capon
goble	2.82	capon
zobel	2.34	watson
watson	IT212	barringer
woods	IT204	barringer
capon	A14	watson
lindsey	2.10	woods
barringer	2.125	<i>null</i>

lecturer	roomno	appraiser	aproom
kahn	IT206	watson	IT212
bush	2.26	capon	A14
goble	2.82	capon	A14
zobel	2.34	watson	IT212
watson	IT212	barringer	2.125
woods	IT204	barringer	2.125
capon	A14	watson	IT212
lindsey	2.10	woods	IT204

37

Exercises

Consider the University db with the tables:

student(studno,name,hons,tutor,year)

staff(lecturer,roomno)

enrolled(studno,courseno,labmark,exammark)

Write queries in relational algebra that return the following:

1. The numbers of courses where a student had a better exam mark than lab mark.
2. The names of the lecturers who are tutoring a student who had an exam mark worse than the lab mark.
3. The names of the lecturers who are tutoring a 3rd year student.
4. The room numbers of the lecturers who are tutoring a 3rd year student.
5. The names of the lecturers who are tutoring more than one student
6. The names of the lecturers who are tutoring no more than one student (?!)

38

Exercise: Cardinalities

Consider relations R and S. Suppose \mathbf{X} is a set of attributes of R. What is the minimal and maximal cardinality of the following relations, expressed in cardinalities of R and S?

- $\sigma_{\mathbf{C}}(R)$
- $\pi_{\mathbf{X}}(R)$
- $\pi_{\mathbf{X}}(R)$, if \mathbf{X} is a superkey of R
- $R \times S$
- $R \times S$, if both R and S are nonempty

39

Cardinalities (cntd.)

Suppose Z is the set of common attributes of R and S.

What is the minimal and maximal cardinality of the following relations, expressed in cardinalities of R and S?

- $R \bowtie S$
- $R \bowtie S$, if Z is a superkey of R
- $R \bowtie S$, if Z is the primary key of R,
and there is a foreign key constraint
 $S(Z) \text{ REFERENCES } R(Z)$

40

Duplicate Elimination

Real DBMSs implement a version of relational algebra that operates on multisets (“bags”) instead of sets.

(Which of these operators may return bags, even if the input consists of sets?)

For the bag version of relational algebra, there exists a duplicate elimination operator δ .

$$\text{If } R = \begin{array}{|c|c|} \hline A & B \\ \hline 1 & 2 \\ \hline 3 & 4 \\ \hline 3 & 4 \\ \hline 1 & 2 \\ \hline \end{array}, \text{ then } \delta(R) = \begin{array}{|c|c|} \hline A & B \\ \hline 1 & 2 \\ \hline 3 & 4 \\ \hline \end{array}$$

41

Aggregation

- Often, we want to retrieve aggregate values, like the “sum of salaries” of employees, or the “average age” of students.
- This is achieved using aggregation functions, such as SUM, AVG, MIN, MAX, or COUNT.
- Such functions are applied by the grouping and aggregation operator γ .

$$\text{If } R = \begin{array}{|c|c|} \hline A & B \\ \hline 1 & 2 \\ \hline 3 & 4 \\ \hline 3 & 5 \\ \hline 1 & 1 \\ \hline \end{array}, \text{ then } \gamma_{\text{SUM}(A)}(R) = \begin{array}{|c|} \hline \text{SUM}(A) \\ \hline 8 \\ \hline \end{array}$$
$$\text{and } \gamma_{\text{AVG}(B)}(R) = \begin{array}{|c|} \hline \text{AVG}(B) \\ \hline 3 \\ \hline \end{array}$$

42

Grouping and Aggregation

- More often, we want to retrieve aggregate values for groups, like the “sum of employee salaries” per department, or the “average student age” per faculty.
- As additional parameters, we give γ attributes that specify the criteria according to which the tuples of the argument are grouped.
- E.g., the operator $\gamma_{A, \text{SUM}(B)}(R)$
 - partitions the tuples of R in groups that agree on A ,
 - returns the sum of all B values for each group.

If $R =$

A	B
1	2
3	4
3	5
1	3

, then $\gamma_{A, \text{SUM}(B)}(R) =$

A	SUM(B)
1	5
3	9

43

Exercise: Identities

Consider relations $R(a,b)$, $R1(a,b)$, $R2(a,b)$, and $S(c,d)$.

For each identity below, find out whether or not it holds for all possible instances of the relations above.


1. $\sigma_{d>5}(R \bowtie_{a=c} S) = R \bowtie_{a=c} \sigma_{d>5}(S)$
2. $\pi_a(R1) \cap \pi_a(R2) = \pi_a(R1 \cap R2)$
3. $(R1 \cup R2) \bowtie_{a=c} S = (R1 \bowtie_{a=c} S) \cup (R2 \bowtie_{a=c} S)$
4. $\pi_c(\sigma_{d>5}(S)) = \sigma_{d>5}(\pi_c(S))$

- If an identity holds, provide an argument why this is the case.
- If an identity does not hold, provide a counterexample, consisting of an instance of the relations concerned and an explanation why the two expressions have different values for that instance.

44

Join: An Observation

Employee	Department
Brown	A
Jones	B
Smith	B



Department	Head
B	Black
C	White

Employee	Department	Head
Jones	B	Black
Smith	B	Black

Some tuples don't contribute to the result, they get lost.

45

Outer Join

- An outer join extends those tuples with null values that would get lost by an (inner) join.
- The outer join comes in three versions
 - left: keeps the tuples of the left argument, extending them with nulls if necessary
 - right: ... of the right argument ...
 - full: ... of both arguments ...

46

(Natural) Left Outer Join

Employee

Employee	Department
Brown	A
Jones	B
Smith	B

Department

Department	Head
B	Black
C	White

Employee ⋈^{Left} Department

Employee	Department	Head
Brown	A	null
Jones	B	Black
Smith	B	Black

(Natural) Right Outer Join

Employee

Employee	Department
Brown	A
Jones	B
Smith	B

Department

Department	Head
B	Black
C	White

Employee ⋈^{Right} Department

Employee	Department	Head
Jones	B	Black
Smith	B	Black
null	C	White

(Natural) Full Outer Join

Employee

Employee	Department
Brown	A
Jones	B
Smith	B

Department

Department	Head
B	Black
C	White

Employee \bowtie^{Full} Department

Employee	Department	Head
Brown	A	null
Jones	B	Black
Smith	B	Black
null	C	White

49

Summary

- Relational algebra is a query language for the relational data model
- Expressions are built up from relations and unary and binary operators
- Operators can be divided into set theoretic, renaming, removal and combination operators (plus extended operators)
- Relational algebra is the target language into which user queries are translated by the DBMS
- Identities allow one to rewrite expressions into equivalent ones, which may be more efficiently executable (\rightarrow query optimization)

50

References

In preparing these slides I have used several sources.
The main ones are the following:

Books:

- A First Course in Database Systems, by J. Ullman and J. Widom
- Fundamentals of Database Systems, by R. Elmasri and S. Navathe

Slides from Database courses held by the following people:

- Enrico Franconi (Free University of Bozen-Bolzano)
- Carol Goble and Ian Horrocks (University of Manchester)
- Diego Calvanese (Free University of Bozen-Bolzano) and Maurizio Lenzerini (University of Rome, "La Sapienza")

In particular, a number of figures are taken from their slides.